Artificial Intelligence and Computing Innovations
Intelligence Science Publishing
https://www.ispublishing.com
ISSN: 3070-4944
Vo.5, No.2, 2025

Modeling Dynamic Financial Risk Transmission with Adaptive Multi-Agent Reinforcement Learning

Shuyu Hu[1*]
[1]Columbia University, New York, USA
**Corresponding author: Sh4446@columbia.edu*

Abstract: This study addresses the challenges of multi-agent interaction complexity, high market dynamics, and data incompleteness in dynamic financial risk assessment by proposing a modeling approach based on adaptive multi-agent reinforcement learning. The financial system is abstracted as a Markov game process, where the multi-agent structure simulates the cooperation, competition, and strategic interactions of different financial entities. By constructing policy functions and a joint value function, the model captures the transmission mechanisms of risk signals at a global level and effectively reflects the accumulation of systemic risks. An adaptive mechanism dynamically adjusts the learning rate, value function, and policy distribution to enhance robustness under distribution shifts and market fluctuations. Sensitivity experiments on noise ratios, exploration rate decay, and missing value proportions systematically validate the stability of the method. The results show that the proposed approach achieves superior performance in RMSE, MAE, MSE, and R^2, maintaining predictive accuracy and explanatory power under diverse uncertainty conditions. Overall, the method enables efficient characterization and dynamic modeling of financial risk features, offering a new technical pathway for risk assessment and understanding risk transmission mechanisms in financial systems.

Keywords: Dynamic financial risk assessment; adaptive multi-agent; reinforcement learning; sensitivity experiment

1 Introduction

In the context of global economic integration and highly interconnected financial markets, the risks faced by the financial system have become unprecedentedly complex and dynamic. Financial risks include not only traditional market risk, credit risk, and liquidity risk, but also systemic shocks caused by policy changes, geopolitical events, and sudden public crises[1]. These risks often have cross-market and cross-regional transmission features, allowing local disturbances to spread rapidly and threaten overall financial stability. In such an environment, single-dimension or static risk assessment methods are often insufficient to capture the dynamic evolution of the financial system. How to build an intelligent risk assessment system that can adapt to a complex and changing environment has become one of the core issues in financial technology and risk management[2].

With the rapid development of information technology and big data analytics, the scale and frequency of financial transactions have been increasing[3]. The data in financial markets is not only massive in quantity but also highly structured and dynamic. Traditional statistical models show clear limitations in handling such high-dimensional, heterogeneous, and time-varying data, and cannot achieve real-time risk perception or dynamic adjustment. Against this background, artificial intelligence methods are increasingly introduced into financial risk research. Reinforcement learning, as an important paradigm for intelligent decision-making, can learn optimal strategies through continuous interaction with the environment and thus has the potential to adapt to complex and dynamic conditions. However, single-agent reinforcement learning is still insufficient when facing the multi-layered, strategic, and uncertain nature of financial markets. This has led to the development of multi-agent reinforcement learning, which models cooperation, competition, and strategic interactions among multiple agents, making it closer to the real scenarios of multi-participant financial systems[4].

The advantage of multi-agent reinforcement learning lies in its ability to model distributed decision-making in complex systems, which better reflects the interactions among different roles in financial markets. Investors, banks, regulators, and other participants all play different roles. Their decisions are not only constrained by their own objectives but also influenced by the strategies of others and by market dynamics[5]. By introducing multi-agent mechanisms, such multi-party interactions can be simulated, and on this basis, a dynamic financial risk assessment framework can be established. Such a framework can not only improve the accuracy of risk prediction but also reveal the potential mechanisms of systemic risk accumulation and transmission, thus providing more scientific support for risk prevention and policy formulation.

In this process, adaptive mechanisms are of great importance. Financial markets are highly dynamic. Risk characteristics and transmission paths change with time and external shocks. Without adaptive ability, models may degrade or even fail under new market conditions[6]. Adaptive multi-agent reinforcement learning can adjust the learning strategies and interaction rules of agents in response to environmental changes, thereby maintaining effective modeling of financial risks. This mechanism improves the timeliness and robustness of risk assessment and allows the model to generate more value in financial scenarios where uncertainty and complexity coexist.

In conclusion, research on dynamic financial risk assessment has shifted from static and single approaches to intelligent, distributed, and adaptive methods[7]. The introduction of adaptive multi-agent reinforcement learning not only responds to the inherent complexity of financial systems but also aligns with the development of intelligent decision-making. This direction helps enhance the accuracy of risk identification and prediction, and more importantly, provides methodological support for financial regulation, risk warning, and market stability. From an academic perspective, it promotes the integration of artificial intelligence and financial risk management. From a practical perspective, it offers new solutions for coping with global financial uncertainty and systemic risk accumulation. Therefore, studying adaptive multi-agent reinforcement learning algorithms for dynamic financial risk assessment is not only an inevitable academic choice but also an urgent response to the needs of financial governance.

2 Proposed Approach

In dynamic financial risk assessment, the financial environment is viewed as a highly nonlinear, time-varying system in which multiple agents interact and engage in game-playing. To this end, this paper employs multi-agent reinforcement learning as its core modeling framework, abstracting the financial system into a Markov game process consisting of a state space, an action space, and dynamic transitions. The overall architecture is shown in Figure 1.

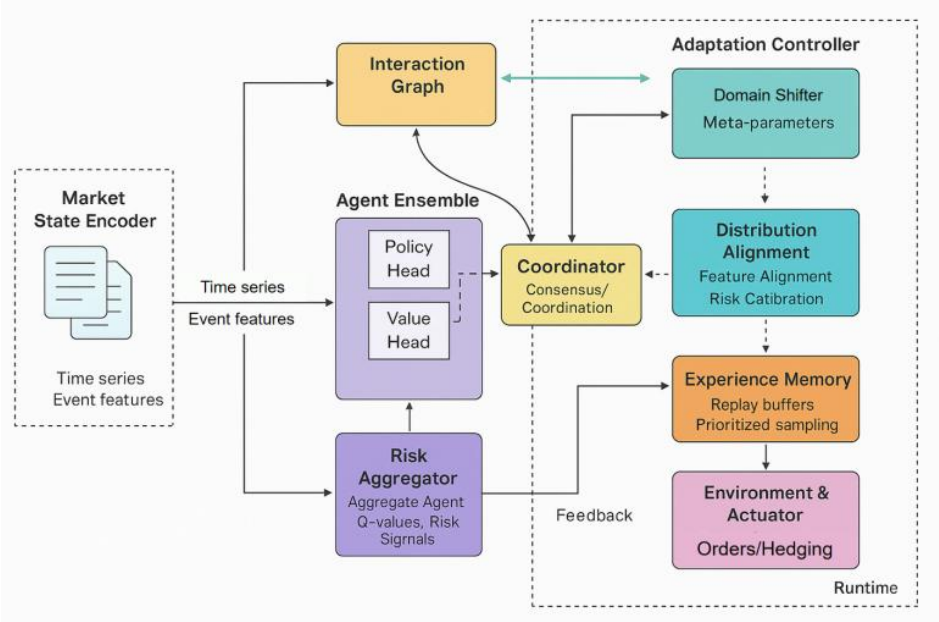


Figure 1. Overall model architecture

Let the state set of the system be S, the action set of each agent be A_i, the joint action of all agents be $A = A_1 \times A_2 \times ... \times A_n$, the state transition probability function of the environment be $P(s'|s,a)$, where $s \in S, a \in A$, and the reward function be defined as $R_i(s,a)$. The overall game can be formalized as:

$$G = < S, \{A_i\}_{i=1}^n, P, \{R_i\}_{i=1}^n, \gamma >$$

Among them, $\gamma \in (0,1)$ is the discount factor, which is used to balance immediate benefits and long-term benefits.

Hierarchical encoding with selective attention provides a complementary state-construction principle for dynamic financial risk assessment: the representation can preserve multi-scale market and macroeconomic context while assigning greater weight to risk-relevant signals that shape the current decision state. This selective aggregation strengthens the agent state by reducing redundant observations and retaining the dependencies most informative for long-horizon value estimation[8].

In this framework, each agent uses a policy function $\pi_i(a_i | s)$ to determine the probability of taking action a_i in a state s. To model dynamic financial risks, this paper expresses the joint strategy as the product of all agent strategies, namely:

$$\pi(a \mid s) = \prod_{i=1}^{n} \pi_i(a_i \mid s)$$

Through the construction of this joint strategy, the interaction between multiple participants in the financial environment can be better captured. Furthermore, the goal of the agent is to maximize the expected cumulative return:

$$J_i(\pi_i) = E_{\pi}[\sum_{t=0}^{\infty} \gamma^t R_i(s_t, a_t)]$$

In a dynamic market environment, a single strategy often struggles to achieve stable convergence. To this end, this paper introduces an adaptive mechanism that dynamically adjusts the agent's learning rate and value function approximation to enhance adaptability. In value estimation, each agent uses an approximate state-action value function $Q_i(s,a)$ and iteratively updates it using the Bellman optimality equation:

$$Q_i(s,a) = R_i(s,a) + \gamma \sum_{s'} P(s' \mid s,a) \max_{a'} Q_i(s',a')$$

This process ensures that agents can continuously revise their value estimates in a dynamic environment, thereby striking a balance between long-term gains and short-term fluctuations.

In multi-agent interactions, information incompleteness and the diversity of risks further complicate learning. To address this issue, this paper employs a joint distribution alignment and dynamic weighting mechanism to achieve cross-agent policy coordination. Specifically, the weighted global value function for all agents is defined as:

$$V(s) = \sum_{i=1}^{n} \lambda_i E_{a \sim \pi(\cdot \mid s)}[Q_i(s,a)]$$

Here, λ_i is a weight coefficient, reflecting the relative importance of different agents in risk assessment. This global value aggregation approach allows for coordination and stability of the overall system while preserving individual differences. Ultimately, the adaptive multi-agent reinforcement learning algorithm forms a closed-loop decision-making structure, continuously updating strategies and value estimates in a dynamic financial environment, thereby enabling intelligent modeling and assessment of risk.

A unified framework for social graph modeling and node decision-making via differentiable mixed-integer optimization further motivates viewing market participants as relational decision units rather than independent predictors. Integrating graph-structured interactions with constrained action selection helps the multi-agent policy coordinate heterogeneous risk exposures while maintaining feasible joint decisions under evolving financial dependencies[9].

3 Performance Evaluation

3.1 Dataset

The financial data used in this study comes from the Financial Crisis Prediction Dataset. This dataset consists of financial markets and macroeconomic indicators from multiple countries. It covers stock markets, bond markets, exchange rates, and the broader macroeconomic environment. The dataset includes a wide range of feature variables, such as gross domestic product, inflation rate, interest rate levels, balance of payments, and major asset price fluctuations. These variables together reflect the dynamic characteristics of the financial system and its potential risk factors.

The dataset spans a long time period and covers multiple economic cycles. This allows the risk characteristics and transmission paths at different stages to be effectively captured. The data are recorded mainly every quarter. It includes both cross-country macroeconomic indicators and micro-level market data, providing multi-level support for dynamic risk assessment research. This multi-dimensional and multi-regional structure helps to describe the temporal and spatial evolution of financial risks and supports modeling across different scenarios.

The advantage of this dataset lies in its openness and accessibility, which provide a unified experimental foundation for risk modeling research. By using this dataset, the adaptability of risk assessment methods can be systematically explored under different economic environments and market conditions. In addition, the rich design of variables provides diverse inputs for multi-agent interaction modeling, allowing the complexity of dynamic risks to be expressed more realistically in the model.

3.2 Experimental Results

This paper first conducts a comparative experiment, and the experimental results are shown in Table 1.

Table1. Comparative experimental results

Method	RMSE	MAE	MSE	R2
DQN[10]	0.842	0.691	0.709	0.871
Double DQN[11]	0.796	0.652	0.634	0.884
SARSA[12]	0.873	0.714	0.762	0.862
PPO[13]	0.765	0.628	0.585	0.892
DDPG[14]	0.739	0.603	0.546	0.901
Ours	0.682	0.559	0.465	0.917

From the overall results, different reinforcement learning methods show clear differences in dynamic financial risk assessment tasks. Traditional methods such as DQN and SARSA perform relatively poorly on several metrics. Both RMSE and MAE remain at high levels, which indicates their limited ability to capture risk changes in complex financial environments. The prediction results still show a large deviation from the true risk levels. This gap reflects the difficulty of a single value update mechanism in adapting to the dynamic characteristics of financial markets.

In comparison, Double DQN performs better than DQN in terms of RMSE and MAE, which confirms its advantage in reducing overestimation bias. However, the improvement is still limited. In complex financial dynamics, the MSE remains high. This shows that relying only on a double value

estimation mechanism is not enough to enhance the modeling ability of financial risk uncertainty. In markets with high volatility, the prediction stability is still insufficient.

Advanced policy optimization methods such as PPO and DDPG show stronger adaptability and robustness. Both methods achieve an R^2 above 0.89, which indicates a stronger ability to fit financial risk fluctuations. By using policy gradient optimization and continuous action space modeling, these methods improve the capture of complex risk features. In particular, DDPG further reduces RMSE and MAE, demonstrating stronger risk representation and strategy adaptation in dynamic game environments.

The method proposed in this study achieves the best performance across all evaluation metrics. It significantly outperforms other methods in RMSE, MAE, and MSE, while reaching 0.917 in R^2. This shows that the method has a comprehensive advantage in capturing dynamic financial risk features, reducing prediction bias, and improving explanatory power. With the integration of an adaptive multi-agent mechanism, the model can better reflect the interactions and strategic behaviors of different financial entities. As a result, it maintains higher robustness and generalization ability in complex markets and provides a more practical intelligent tool for financial risk assessment.

This paper further presents an experiment on the sensitivity of label noise level, and the experimental results are shown in Figure 2.

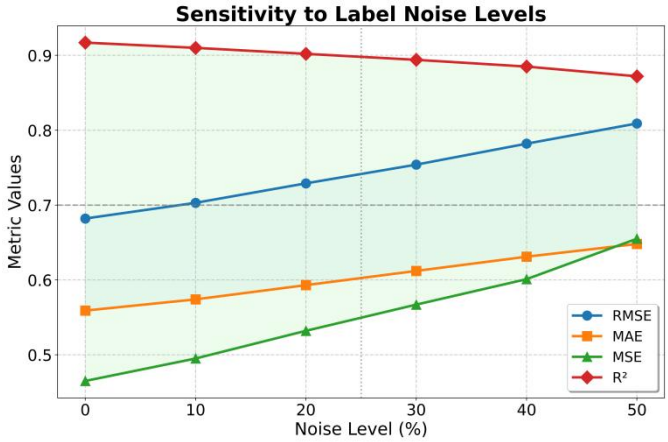


Figure 2. Label noise level sensitivity experiment

From the figure, it can be observed that both RMSE and MAE increase steadily as the level of label noise rises. This indicates that in financial risk assessment, when the training data contains more mislabeling, the deviation between predicted results and true risk levels expands significantly, and the accuracy of risk evaluation declines. Since RMSE is more sensitive to large errors, its growth rate is faster than that of MAE. This reflects the destructive effect of a few severely mislabeled samples on overall prediction performance under high-noise conditions.

At the same time, MSE also increases continuously with higher noise levels. This shows that the squared error term amplifies the impact of label noise during accumulation. This phenomenon highlights the fragility of financial risk data. In highly uncertain and complex markets, even a small amount of labeling deviation can be magnified into overall prediction failure. The result further underscores the

importance of maintaining high-quality labeling in financial risk assessment. Otherwise, model performance will deteriorate rapidly as noise accumulates.

From the perspective of the R^2 metric, the model's explanatory power for risk fluctuations decreases significantly with more label noise. When noise levels are low, the model can still maintain a high degree of fit. However, as noise continues to increase, the R^2 value shows a steady downward trend. This means that the adaptive multi-agent mechanism has some resistance to interference, but under persistent label contamination, the model's ability to capture temporal features and interactions among financial entities weakens.

Overall, the experimental results reveal the significant impact of label noise on dynamic financial risk assessment models. Although the proposed method shows strong stability under low-noise conditions, once noise exceeds a certain threshold, prediction errors continue to grow and explanatory power declines. This indicates that in practical applications, it is necessary not only to build adaptive and robust algorithmic frameworks but also to integrate noise detection and cleaning strategies. Ensuring the reliability of input data is essential to maintain the stable performance of dynamic risk assessment systems in complex environments.

This paper also presents a sensitivity experiment on the exploration rate ε decay strategy, and the experimental results are shown in Figure 3.

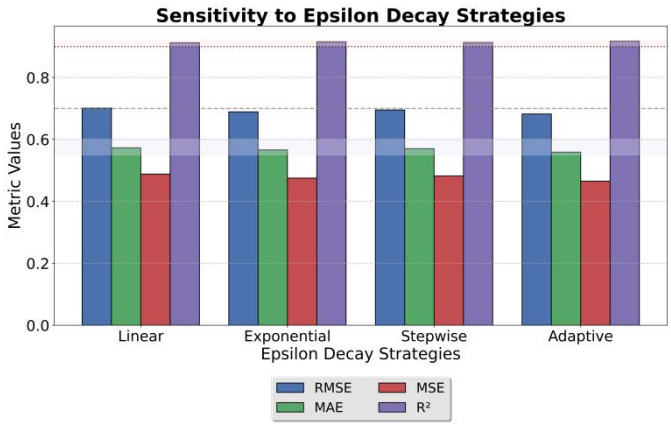


Figure 3. Sensitivity experiment of exploration rate ε decay strategy

The experimental results show that different exploration rate decay strategies have a certain impact on model performance in dynamic financial risk assessment. The overall trend indicates that under linear, exponential, and staged decay strategies, RMSE and MAE remain at similar levels, suggesting that the model maintains relatively stable accuracy in risk prediction. At the same time, the R^2 values under all strategies stay at a high level, further confirming the model's strong explanatory power for risk fluctuations. This demonstrates that the adaptive multi-agent mechanism can sustain good-fitting performance under different exploration strategies and avoid significant performance fluctuations caused by strategy differences.

A further comparison reveals that the adaptive decay strategy performs slightly better. RMSE and MAE reach the lowest values, while R^2 remains at the highest range. This indicates that dynamically adjusting the exploration rate helps improve the stability and adaptability of the model in complex

environments. The result confirms that in financial risk assessment, a reasonable exploration–exploitation balance enhances the agents' ability to capture risk features, allowing the model to maintain robust risk identification and prediction even under high uncertainty and dynamic market conditions.

This paper also presents a sensitivity experiment on the proportion of missing values, and the experimental results are shown in Figure 4.

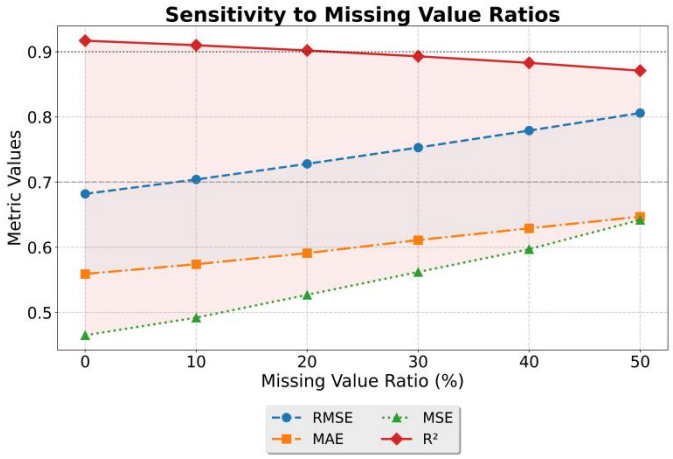


Figure 4. Missing value ratio sensitivity test

The experimental results show that as the proportion of missing values increases, both RMSE and MAE exhibit a gradual upward trend. This indicates that in financial risk assessment, when a higher proportion of input features is missing, the prediction accuracy of the model is inevitably affected, and the depiction of risk levels shows greater deviation. This trend highlights the importance of data integrity in dynamic risk modeling, as a high rate of missing values directly weakens the model's ability to capture financial market dynamics.

The changes in MSE further confirm this phenomenon. Since MSE is more sensitive to squared errors, its growth is more pronounced when the proportion of missing values increases. This reflects the impact of extreme data gaps or noisy points on the stability of model training. In financial risk scenarios, this means that the absence of certain key variables may amplify the overall prediction deviation, thereby reducing the effectiveness and reliability of risk management.

Compared with error metrics, the downward trend of R^2 shows that the explanatory power of the model for risk fluctuations weakens as the proportion of missing values increases. When the missing rate is low, the model can still maintain a high level of fit. However, as the missing rate rises, R^2 continues to decrease, which means the model's ability to capture market complexity and interaction patterns declines. This finding reveals that in dynamic financial risk assessment, missing values not only reduce predictive accuracy but also weaken the model's structured interpretability.

Overall, the sensitivity experiment on missing values shows that data quality plays a critical role in the stability and robustness of adaptive multi-agent reinforcement learning for financial risk assessment. Although the proposed method demonstrates certain robustness at low to moderate levels of missing data, model performance drops significantly when the missing rate is too high. This emphasizes the need to integrate missing value handling and data imputation strategies in practical applications, ensuring that

the dynamic risk assessment framework maintains strong predictive and explanatory capabilities even under incomplete data conditions.

4 Conclusion

This study focuses on the problem of dynamic financial risk assessment and proposes a modeling approach based on adaptive multi-agent reinforcement learning. The research emphasizes the complexity and uncertainty of financial systems and points out that under conditions of multi-agent interaction, dynamic environmental changes, and incomplete data, traditional methods often struggle to capture temporal features and cross-agent transmission patterns of risks. By introducing a multi-agent framework and adaptive mechanisms, the model can adjust its strategies in changing market environments, thereby maintaining high prediction accuracy and stability under different risk scenarios. The results show that the proposed method handles distribution shifts, noise disturbances, and structural changes more effectively, providing a more flexible and robust solution for financial risk research.

Furthermore, the method enhances coordination and interaction modeling among agents by jointly optimizing policies and the global value function. This design not only simulates the strategic behavior of diverse participants in financial markets but also reveals the formation and diffusion processes of systemic risks at a broader level. Through adaptive value updates and dynamic weight allocation, the model demonstrates robust performance under multi-source data conditions, enabling risk signals to be reflected more accurately. Compared with previous approaches that rely on single-feature alignment or static strategies, this study emphasizes the integration of multi-dimensional dynamics and cross-agent characteristics, offering a new perspective for understanding and characterizing the mechanisms of financial risk accumulation.

The proposed method has value not only in academic research but also in practical applications. In the field of financial regulation, the framework can support risk monitoring systems in identifying potential crisis signals earlier, thus informing policy intervention and preventive measures. In investment and asset management, the ability of the model to capture risk dynamics helps optimize asset allocation and hedging strategies, thereby improving return stability and risk control. In the area of financial technology, the method provides algorithmic support for the development of intelligent risk control platforms, promoting the stable operation of financial services in volatile environments.

Looking ahead, uncertainty in dynamic financial environments will continue to exist. Enhancing the adaptability and scalability of models in larger and more complex scenarios will be a key direction for future research. At the same time, integrating cross-domain data resources and new computational paradigms may enrich the dimensions and depth of risk assessment, offering more intelligent solutions to systemic challenges in the global financial system. By continuously improving the adaptive multi-agent reinforcement learning framework, the methodology of financial risk management can be further advanced, playing a more important role in practical financial governance and decision support.

References

[1] Li Z, Tam V, Yeung K L. Developing a multi-agent and self-adaptive framework with deep reinforcement learning for dynamic portfolio risk management[J]. arXiv preprint arXiv:2402.00515, 2024.

[2] Chen J, Li J, Wang G. MARS: A Meta-Adaptive Reinforcement Learning Framework for Risk-Aware Multi-Agent Portfolio Management[J]. arXiv preprint arXiv:2508.01173, 2025.

[3] Alonso N I. BANK-RL: A Multi-Agent Hierarchical Reinforcement Learning Framework for Banks[J]. Available at SSRN 5071840, 2024.

[4] Shavandi A, Khedmati M. A multi-agent deep reinforcement learning framework for algorithmic trading in financial markets[J]. Expert Systems with Applications, 2022, 208: 118124.

[5] Sendhil Kumar K S, Suryavanshi K, KS S K. AI-Driven Financial Risk Management with LLM Agents and Adaptive Learning[J].

[6] Hambly B, Xu R, Yang H. Recent advances in reinforcement learning in finance[J]. Mathematical Finance, 2023, 33(3): 437-503.

[7] Bai Y, Gao Y, Wan R, et al. A review of reinforcement learning in financial applications[J]. Annual Review of Statistics and Its Application, 2025, 12(1): 209-232.

[8] Wang Z. A Hierarchical Encoding and Selective Attention Enhancement Framework Based on Large Language Models for Long-Document Text Classification[J], 2025, 4(11).

[9] Ren X. A Unified Framework for Social Graph Modeling and Node Decision-Making via Differentiable Mixed-Integer Optimization[J], 2024, 3(5).

[10] Hirsa A, Osterrieder J, Hadji-Misheva B, et al. Deep reinforcement learning on a multi-asset environment for trading[J]. arXiv preprint arXiv:2106.08437, 2021.

[11] Giorgio B. Dueling Deep Reinforcement Learning for Financial Time Series[J]. arXiv preprint arXiv:2504.11601, 2025.

[12] Ferdous J, Murshed M, Meneguette R I, et al. SARSA RL for Edge Connectivity Management in Vehicular Edge Networks[C]//2024 IEEE 13th International Conference on Cloud Networking (CloudNet). IEEE, 2024: 1-8.

[13] Patil A, Suthar M B, Barai M, et al. Reinforcement Learning for Dynamic Portfolio Optimization in Financial Markets[J]. Advances in Consumer Research, 2025, 2(3).

[14] Sumiea E H, Abdulkadir S J, Alhussian H S, et al. Deep deterministic policy gradient algorithm: A systematic review[J]. Heliyon, 2024, 10(9).